Explicit Lower Bounds via Geometric Complexity Theory

[Extended Abstract]

Peter Bürgisser*
Institute of Mathematics
University of Paderborn
D-33098 Paderborn, Germany
pbuerg@math.upb.de

Christian Ikenmeyer†
Institute of Mathematics
University of Paderborn
D-33098 Paderborn, Germany
ciken@math.upb.de

ABSTRACT

We prove the lower bound $\underline{R}(\mathcal{M}_m) \geq \frac{3}{2}m^2 - 2$ on the border rank of $m \times m$ matrix multiplication by exhibiting explicit representation theoretic (occurence) obstructions in the sense of Mulmuley and Sohoni's geometric complexity theory (GCT) program. While this bound is weaker than the one recently obtained by Landsberg and Ottaviani, these are the first significant lower bounds obtained within the GCT program. Behind the proof is the new combinatorial concept of *obstruction designs*, which encode highest weight vectors in $\mathsf{Sym}^d \bigotimes^3 (\mathbb{C}^n)^*$ and provide new insights into Kronecker coefficients.

Categories and Subject Descriptors

F.1.3 [**Computation by abstract devices**]: Complexity Measures and Classes; F.2.1 [**Analysis of Algorithms and Problem Complexity**]: Numerical Algorithms and Problems—*Computations on polynomials*

General Terms

Algorithms, Theory

Keywords

geometric complexity theory, tensor rank, matrix multiplication, Kronecker coefficients, permanent versus determinant

Acknowledgments

We thank Matthias Christandl, Jon Hauenstein, and J.M. Landsberg for important discussions.

1. INTRODUCTION

The complexity of matrix multiplication is captured by the rank of the matrix multiplication tensor, a quantity that, despite intense research efforts, is little understood. Strassen [21] already observed that the closely related notion of border rank has a natural formulation as a specific orbit closure problem. The work [3] applied and further developed the collection of ideas from Mulmuley and Sohoni [16, 17] to the tensor framework, which is simpler than the one for permanent versus determinant. However, the lower bound obtained in [3] for the border rank $\underline{R}(\mathcal{M}_m)$ of the $m \times m$ matrix multiplication tensor $\mathcal{M}_m$ is ridiculously small. In this work, we considerably improve this bound by explicitly constructing representation theoretic obstructions in a combinatorial way, cf. Thm. 6.2. While our bound obtained is slightly below the one by Strassen [20] and Lickteig [15], and also weaker than the very recent improvement by Landsberg and Ottaviani [13], we think that it is still worthwhile to publish our method now. The reason is that we show for the first time that significant lower bounds can be obtained with geometric complexity theory (GCT). Our actual construction of highest weight vectors is combinatorial (e.g., the vanishing of the determinants of the matrices involved is enforced by requiring that the matrix has duplicate columns). So there is a lot of leeway left for improvements. As further evidence for this, we note that recently, in collaboration with Jon Hauenstein and J.M. Landsberg, we managed to prove $\underline{R}(\mathcal{M}_2) = 7$ using an explicit construction of highest weight vectors of weight $\lambda = (5, 5, 5, 5)^3$ and relying on computer calculations. This is remarkable, since this was a long-standing open problem since the 70s, which was only settled in 2005 by Landsberg [11] using very different methods.

As a further contribution, we clarify the discussion on the feasibility of the GCT approach by pointing out that in a modification of the approach, proving lower bounds is actually *equivalent* to providing the existence of obstructions (in the sense of highest weight vectors instead of just highest weights), cf. Prop. 3.3. Finally, we obtain a new characterization of Kronecker coefficients (Thm. 5.7) with interesting consequences (Cor. 5.8, Prop. 5.9).

This extended abstract contains results from the PhD thesis of the second author [9]. Several proofs had to be omitted for lack of space.

2. ORBIT CLOSURE PROBLEMS

2.1 Border Rank

Consider $W := \bigotimes^3 \mathbb{C}^m$. The *rank* $R(w)$ of a tensor $w \in W$ is defined as the minimum $r \in \mathbb{N}$ such that w can be written as a sum of r tensors of the form $w_1 \otimes w_2 \otimes w_3$ with $w_i \in \mathbb{C}^m$.

*partially supported by DFG-grant BU 1371/3-2
†partially supported by DFG-grant BU 1371/3-2

MSC class: 68Q17, 20C30, 14L24

Strassen proved [19] that, up to a factor of two, $R(w)$ equals the minimum number of nonscalar multiplications sufficient for evaluating the bilinear map $(\mathbb{C}^m)^* \times (\mathbb{C}^m)^* \to \mathbb{C}^m$ corresponding to w. The *border rank* $\underline{R}(w)$ of a tensor $w \in W$ is defined as the smallest $r \in \mathbb{N}$ such that w can be obtained as the limit of a sequence $w_k \in W$ with tensor rank $R(w_k) \le r$ for all k. Border rank is a natural mathematical notion that has played an important role in the discovery of fast algorithms for matrix multiplication, see [2, Ch. 15].

Now let $n \ge m$ and think of W as embedded in $V := \bigotimes^3 \mathbb{C}^n$. The group $G := \mathsf{GL}_n^3$ acts on V via $(g_1, g_2, g_3)(w_1 \otimes w_2 \otimes w_3) := g_1(w_1) \otimes g_2(w_2) \otimes g_3(w_3)$. We shall denote by $Gw := \{gw \mid g \in G\}$ the *orbit* of v and call its closure $\overline{Gw}$ with respect to the euclidean topology the *orbit closure* of w.

It will be convenient to use Dirac's bra-ket notation. So $(|i\rangle)_{1\le i\le n}$ denotes the standard basis of $\mathbb{C}^n$ and $|ijk\rangle$ is a short hand for $|i\rangle \otimes |j\rangle \otimes |k\rangle \in V$. We call $\mathcal{E}_n := \sum_{i=1}^n |iii\rangle \in V$ the *n-th unit tensor*.

Suppose that $\underline{R}(w) \ge m$ to avoid trivial cases. Then it is easy to see that $\underline{R}(w) \le n$ iff $w \in \overline{G\mathcal{E}_n}$, cf. [21].

The tensor corresponding to the $m \times m$ matrix multiplication map can be succinctly written as

$$\mathcal{M}_m := \sum_{i,j,l=1}^{m} |(i,j)(j,l)(l,i)\rangle \in \bigotimes^3 \mathbb{C}^{m\times m}. \tag{2.1}$$

2.2 Approximate Determinantal Complexity

We switch now the scenario and take $V := \mathsf{Sym}^n \mathbb{C}^{n^2}$, which is the homogeneous part of degree n of the polynomial ring $\mathbb{C}[X_1, \dots, X_{n^2}]$. The determinant $\det_n$ of an $n \times n$ matrix in these variables is an element of V. The group $G := \mathsf{GL}_{n^2}$ acts on V by linear substitution. Further, let $m < n$, and put $z := X_{m^2+1}$, $W := \mathsf{Sym}^m \mathbb{C}^{m^2}$. We define the *determinantal orbit closure complexity* $\mathsf{docc}(f)$ of $f \in W$ as the minimal n such that $z^{n-m} f \in \overline{G\det_n}$.

In [16] Mulmuley and Sohoni conjectured the following:

$$\mathsf{docc}(\mathrm{per}_m) \text{ is not polynomially bounded in } m. \tag{2.2}$$

Here $\mathrm{per}_m \in W$ denotes the permanent of the $m \times m$ matrix in the variables $X_1, \dots, X_{m^2}$.

An affirmative answer to this conjecture implies that $\det_n$ cannot be computed by weakly skew circuits of size polynomial in m, (cf. [4]), which is a version of Valiant's Conjecture [23].

2.3 Unifying Notation

The *tensor scenario* and the *polynomial scenario* discussed before have much in common and we strive to treat both situations simultaneously. Hence for fixed n and m we want to use the following notation summarized in the following table.

notation	determinantal complexity ($n \ge m+1$)	border rank ($n \ge m^2$)
G	GL_{n^2}	$\mathsf{GL}_n \times \mathsf{GL}_n \times \mathsf{GL}_n$
V	$\mathsf{Sym}^n \mathbb{C}^{n^2}$	$\bigotimes^3 \mathbb{C}^n$
$\eta = \dim V$	$\binom{n^2+n-1}{n}$	n^3
$W \subseteq V$	$\mathsf{Sym}^m \mathbb{C}^{m^2}$	$\bigotimes^3 \mathbb{C}^{m^2}$
$\mathfrak{h} := \mathfrak{h}_{m,n} \in W$	$z^{n-m}\mathrm{per}_m$	$\mathcal{M}_m$
$\mathfrak{c} := \mathfrak{c}_n \in V$	$\det_n$	$\mathcal{E}_n$

The point $\mathfrak{h}$ stands for the hard problem for which we want to prove lower bounds and the orbit closure $\overline{G\mathfrak{c}_n}$ is exactly the set of all points with complexity at most n. In both scenarios, for a given m, we try to find n as large as possible such that

$$\mathfrak{h}_{m,n} \notin \overline{G\mathfrak{c}_n}.$$

Since the orbit closure is the smallest closed set containing the orbit, this is equivalent to proving $\overline{G\mathfrak{h}_{m,n}} \not\subseteq \overline{G\mathfrak{c}_n}$. If we want to treat $G\mathfrak{c}$ and $G\mathfrak{h}$ simultaneously, we just write Gv.

3. THE FLIP VIA OBSTRUCTIONS

Let $V \simeq \mathbb{C}^\eta$ and $v \in V$ in one of the two scenarios above. We write $\mathbb{C}[V] := \mathbb{C}[T_1, \dots, T_\eta]$ for the ring of polynomial functions on V. It is a fundamental fact from algebraic geometry that the orbit closures $\overline{Gv}$ (defined via the euclidean topology) are in fact Zariski closed, i.e., zero sets of polynomials on V (cf. [10, AI.7.2]). This immediately implies the following observation.

PROPOSITION 3.1. *Let $\mathfrak{h} \in V$. If $\mathfrak{h} \notin \overline{G\mathfrak{c}}$, then there exists a polynomial $f \in \mathbb{C}[V]$ that vanishes on $\overline{G\mathfrak{c}}$ but not on $\mathfrak{h}$.*

We call such polynomials f that separate $\mathfrak{h}$ from $\overline{G\mathfrak{c}}$ *polynomial obstructions*. By Prop. 3.1, they are guaranteed to exist if $\mathfrak{h} \notin \overline{G\mathfrak{c}}$. We shall address the questions whether there are "short encodings" of polynomial obstructions f and whether there are "short proofs" that f is an obstruction. Representation theory provides a natural framework to address these questions.

3.1 Highest Weight Vectors

We recall some facts from representation theory [8]. Let $\mathscr{V}$ be a rational GL_N-representation. For a given $z \in \mathbb{Z}^N$, a *weight vector $f \in \mathscr{V}$ of weight z* is defined by the following property: $\mathrm{diag}(\alpha) f = \alpha_1^{z_1} \alpha_2^{z_2} \cdots \alpha_N^{z_N} f$ for all $\alpha \in (\mathbb{C}^\times)^N$.

Let $U_N \subseteq \mathsf{GL}_N$ denote the group of upper triangular matrices with 1s on the main diagonal, the so-called *maximal unipotent group*. A weight vector $f \in \mathscr{V}$ that is fixed under the action of U_N, i.e., $\forall u \in U_N : uf = f$, is called a *highest weight vector (HWV) of $\mathscr{V}$*. The vector space of HWVs of weight λ is denoted by $\mathsf{HWV}_\lambda(\mathscr{V})$. The following is well known.

LEMMA 3.2. *Each irreducible GL_N-representation $\mathscr{V}$ contains, up to scalar multiples, exactly one HWV f. The representation $\mathscr{V}$ is the linear span of the GL_N-orbit of f. Two irreducible representations are isomorphic iff the weights of their HWVs coincide.*

The weight $\lambda \in \mathbb{Z}^N$ of a HWV is always nondecreasing. It describes the isomorphy type of $\mathscr{V}$. We denote by $\{\lambda\}$ the irreducible GL_N-representation with highest weight λ, called *Weyl-module*. It is a well known fact that $\mathscr{V}$ splits into a direct sum of irreducible GL_N-representations.

What has been said for GL_N extends in a straightforward way to representations $\mathscr{V}$ of the group $\mathsf{GL}_n \times \mathsf{GL}_n \times \mathsf{GL}_n$. A weight vector $f \in \mathscr{V}$ of weight $z \in \mathbb{Z}^n \times \mathbb{Z}^n \times \mathbb{Z}^n$ satisfies

$$(\mathrm{diag}(\alpha^{(1)}), \mathrm{diag}(\alpha^{(2)}), \mathrm{diag}(\alpha^{(3)}))f = \textstyle\prod_{k=1}^3 \prod_{i=1}^n (\alpha_i^{(k)})^{z_i^{(k)}}$$

for all $\alpha^{(k)} \in (\mathbb{C}^\times)^n$. The type of irreducible $\mathsf{GL}_n \times \mathsf{GL}_n \times \mathsf{GL}_n$-representations is given by triples $\lambda = (\lambda^{(1)}, \lambda^{(2)}, \lambda^{(3)})$, where $\lambda^{(k)}$ is a highest weight for GL_n.

3.2 HWV Obstructions

We return to our two scenarios. The action of the group G on V induces an action of G on $\mathbb{C}[V]$ defined by $(gf)(x) := f(g^{-1}x)$ for $g \in G$, $f \in \mathbb{C}[V]$, $x \in V$. This action respects the degree d part $\mathscr{V} = \mathbb{C}[V]_d$. Let $I(Gv) = I(\overline{Gv})$ denote the vanishing ideal of the orbit Gv and let $\mathbb{C}[\overline{Gv}] := \mathbb{C}[V]/I(\overline{Gv})$ denote the coordinate ring of $\overline{Gv}$.

The following result shows that when searching for polynomial obstructions, we can restrict ourselves to HWVs.

PROPOSITION 3.3. *Let $\mathfrak{h} \in V$. If $\mathfrak{h} \notin \overline{G\mathfrak{c}}$, then there exists some HWV $f_\lambda \in \mathbb{C}[V]$ of some weight λ such that f_λ vanishes on $\overline{G\mathfrak{c}}$, but $f_\lambda(g\mathfrak{h}) \neq 0$ for some $g \in G$.*

PROOF. The fact $f(\overline{G\mathfrak{c}}) = 0$ means that f is contained in the vanishing ideal $I(\overline{G\mathfrak{c}})$. But $I(\overline{G\mathfrak{c}})$ is a graded G-representation. Hence we can write $f = \sum_{d,\lambda} f_{d,\lambda}$, where $f_{d,\lambda} \in I(\overline{G\mathfrak{c}})_d$ are elements from the isotypic component of type λ in the homogeneous part $I(\overline{G\mathfrak{c}})_d$. By Lem. 3.2, it follows that we can write $f_{d,\lambda} = \sum_i g_{d,\lambda,i} f_{d,\lambda,i}$, where $g_{d,\lambda,i} \in G$ and $f_{d,\lambda,i}$ is a HWV in $I(\overline{G\mathfrak{c}})_d$ of weight λ.

Let $g \in G$ with $f(g\mathfrak{h}) \neq 0$. Then $g_{d,\lambda,i} f_{d,\lambda,i}(g\mathfrak{h}) \neq 0$ for some d, λ, i. This means $f_{d,\lambda,i}(g_{d,\lambda,i}^{-1} g\mathfrak{h}) \neq 0$, which proves the proposition. $\square$

We shall call such f_λ a *HWV obstruction* against $\mathfrak{h} \in \overline{G\mathfrak{c}}$. In Sec. 5.2, we will see that some HWVs have a succinct encoding, which is linear in their degree d. It would be interesting to have bounds on d in terms of n.

An *occurence obstruction* against $\mathfrak{h} \in \overline{G\mathfrak{c}}$, as introduced by Mulmuley and Sohoni [17, Def. 1.2], is a highest weight λ for G such that irreducible G-representations of type λ do not occur in $\mathbb{C}[\overline{G\mathfrak{c}}]$, but some irreducible G-representation of type λ does occur in $\mathbb{C}[\overline{G\mathfrak{h}}]$. These properties can be rephrased as follows:

- *All* HWVs in $\mathbb{C}[V]$ of weight λ vanish at $\overline{G\mathfrak{c}}$;
- There exists some HWV f_λ in $\mathbb{C}[V]$ of weight λ that does not vanish on $\overline{G\mathfrak{h}}$.

If λ is an occurence obstruction against $\mathfrak{h} \in \overline{G\mathfrak{c}}$, then there exists a HWV obstruction f_λ of weight λ. But the converse is not true in general, see for instance the discussion on Strassen's invariant in [3]. Clearly, if the irreducible represenation corresponding to λ occurs in $\mathbb{C}[V]$ with high multiplicity, then item one above is much harder to satisfy for occurence obstructions.

While Prop. 3.3 tells us that $\mathfrak{h} \notin \overline{G\mathfrak{c}}$ can, in principle, always be proven by exhibiting a HWV obstruction, it is unclear whether this is also the case for occurence obstructions. We state this as an important open problem.

PROBLEM 3.4. *For the scenarios in Subsec. 2.3, if $\mathfrak{h}_{m,n} \notin \overline{G\mathfrak{c}_n}$, is there an occurence obstruction proving this?*

Mulmuley and Sohoni conjecture that (2.2) can be proved with occurence obstructions, see [17, §3].

4. EXPLICIT HWVS

4.1 Combinatorics

A *partition* λ is a finite sequence of nonincreasing natural numbers. The number of nonzero elements in λ is called its *length* $\ell(\lambda)$. The set of partitions is closed under elementwise addition, written as $\lambda + \mu$. For a fixed $n \in \mathbb{N}$ and a partition λ of length at most n let $\lambda^* := (-\lambda_n, \ldots, -\lambda_1)$ denote the *dual of* λ. We call $|\lambda| := \sum_i \lambda_i$ the *size* of λ. If λ satisfies $|\lambda| = d$ and $\ell(d) \le n$, then we write $\lambda \vdash_{\overline{n}} d$. If we do not specify the size, we just write $\lambda \vdash_{\overline{n}}$, and if we do not specify the length, we write $\lambda \vdash d$.

A pictorial description of partitions is given by *Young diagrams*, which are upper-left-justified arrays having λ_i boxes in the ith row. The partitions $\ell \times k$ correspond to rectangular Young diagrams with ℓ rows and k columns. In particular $(r) = 1 \times r$ stands for the partition with a single row and r boxes. The symbol $c \ulcorner r$ shall denote the *hook partition* $(c \times 1) + (r-1)$. Note that $(c \ulcorner r) \vdash_{\overline{c}} r + c - 1$. When reflecting a Young diagram λ at the diagonal from the upper left to the lower right we get a Young diagram again, which we call the *transposed Young diagram* ${}^t\lambda$. Note that the number of boxes of λ in column i equals ${}^t\lambda_i := ({}^t\lambda)_i$.

For partition triples $\lambda = (\lambda^{(1)}, \lambda^{(2)}, \lambda^{(3)})$ we use the short notation $\lambda^{(1)}, \lambda^{(2)}, \lambda^{(3)} \vdash_{\overline{n_1,n_2,n_3}} d_1, d_2, d_3$, which stands for $\lambda^{(k)} \vdash_{\overline{n_k}} d_k$ for all $k \in \{1,2,3\}$. If $n_1 = n_2 = n_3 =: n$ and $d_1 = d_2 = d_3 =: d$ we use the even shorter syntax $\lambda \vdash^{*}_{\overline{n}} d$.

4.2 Polarization, Restitution, and Projections

For a finite group G let $\mathbb{C}[G]$ denote its group algebra. Let $\mathscr{V}$ be a rational G-representation. We use the short notation $\langle f|p|v\rangle := \sum_{g\in G} \alpha_g \langle f|g|v\rangle$ for $\langle f| \in \mathscr{V}^*$, $|v\rangle \in \mathscr{V}$, $\alpha_g \in \mathbb{C}$, and $p = \sum_{g \in G} \alpha_g g \in \mathbb{C}[G]$. Here $\langle f|g|v\rangle := f(gv)$.

The symmetric group S_D acts on $\bigotimes^D \mathbb{C}^N$ via

$$\pi(v_1 \otimes v_2 \otimes \cdots \otimes v_D) := v_{\pi^{-1}(1)} \otimes \cdots \otimes v_{\pi^{-1}(D)}. \qquad (4.1)$$

The subspace of tensors that are invariant under this action is denoted by $\mathsf{Sym}^D \mathbb{C}^N := (\bigotimes^D \mathbb{C}^N)^{\mathsf{S}_D}$. The projection map $\mathcal{P}_D \colon \bigotimes^D \mathbb{C}^N \twoheadrightarrow \mathsf{Sym}^D \mathbb{C}^N$ is given by

$$\mathcal{P}_D := \frac{1}{D!} \sum_{\pi \in \mathsf{S}_D} \pi \in \mathbb{C}[\mathsf{S}_D]. \qquad (4.2)$$

The following easy and well-known lemma from multilinear algebra (e.g. [6, Ch. 1.2]) builds a link between homogeneous polynomials and multilinear forms.

LEMMA 4.1. *Let $W := \mathbb{C}^N$. Let $\langle f| \in \bigotimes^D W^*$. Then define the homogeneous polynomial $f \in \mathbb{C}[W]_D$ via*

$$\forall w \in W : f(w) = \langle f | w^{\otimes D}\rangle.$$

The map $\phi \colon \bigotimes^D W^* \to \mathbb{C}[W]_D$, $\langle f| \mapsto f$ is linear and induces a linear isomorphism $\phi_0 \colon \mathsf{Sym}^D W^* \xrightarrow{\sim} \mathbb{C}[W]_D$. The map $\psi \circ \phi$ is the projection $\mathcal{P}_D$ defined in (4.2).

We remark that $f(w) = \langle f|\mathcal{P}_D|w^{\otimes D}\rangle = \langle f|w^{\otimes D}\rangle$.

4.3 Schur-Weyl Duality and HWVs

The vector space $\bigotimes^D \mathbb{C}^N$ is a $\mathsf{GL}_N \times \mathsf{S}_D$-representation via the commuting actions of S_D and GL_N, defined for S_D in (4.1), and for GL_N as follows:

$$g|w_1 w_2 \cdots w_D\rangle := g|w_1\rangle \otimes g|w_2\rangle \otimes \cdots \otimes g|w_D\rangle, \quad g \in \mathsf{GL}_N.$$

It follows that S_D stabilizes every highest weight vector space $\mathsf{HWV}_\lambda(\bigotimes^d \mathbb{C}^N)$. For $\lambda \vdash D$ let $[\lambda]$ denote the irreducible S_D-representation corresponding to λ, called *Specht module*. The fundamental *Schur-Weyl duality* states that

$$\bigotimes^D \mathbb{C}^N \simeq \bigoplus_{\lambda \vdash_{\overline{N}} D} \{\lambda\} \otimes [\lambda]$$

as $\mathsf{GL}_N \times \mathsf{S}_D$-representations, cf. [8, Sec. 4.2.4]. To make this explicit, we define for $\ell \le N$ the wedge product

$$\langle\widehat{\ell}| := \langle 1| \wedge \langle 2| \wedge \cdots \wedge \langle \ell| := \sum_{\pi \in \mathsf{S}_\ell} \mathrm{sgn}(\pi)\langle 123\cdots\ell|\pi,$$

sometimes called the *Slater-determinant*. Let us point out that the scalar product $\langle\widehat{\ell}|x_1 \otimes x_2 \otimes \cdots \otimes x_\ell\rangle$, where all $x_i \in \mathbb{C}^N$, is just the determinant of the $\ell \times \ell$ matrix

$$\begin{pmatrix} \langle 1|x_1\rangle & \cdots & \langle 1|x_\ell\rangle \\ \vdots & \ddots & \vdots \\ \langle \ell|x_1\rangle & \cdots & \langle \ell|x_\ell\rangle \end{pmatrix}. \tag{4.3}$$

It is easy to verify that $\langle\widehat{\ell}|$ is a HWV of weight $\ell\times(-1)$. Let $\mu \vdash_N D$ be a partition with column lengths ${}^t\mu_1, \ldots, {}^t\mu_{\mu_1}$. Then define

$$\langle\widehat{\mu}| := \bigotimes_{i=1}^{\mu_1} \langle\widehat{{}^t\mu_i}| \in \bigotimes^D (\mathbb{C}^N)^*. \tag{4.4}$$

It is readily checked that $\langle\widehat{\mu}|$ is a HWV of weight μ^*.

Claim 4.2. *Let $\mu \vdash_N D$. The tensors $\langle\widehat{\mu}|\pi$ with $\pi \in \mathsf{S}_D$ generate the vector space $\mathsf{HWV}_{\mu^*}(\bigotimes^D(\mathbb{C}^N)^*)$.*

Proof. Put $W := (\mathbb{C}^N)^*$. Schur-Weyl duality and $[\mu] \simeq [\mu]^*$ implies $\mathsf{HWV}_{\mu^*}\left(\bigotimes^D W\right) \simeq \mathsf{HWV}_{\mu^*}(\{\mu^*\}) \otimes [\mu]$. But $\mathsf{HWV}_{\mu^*}(\{\mu^*\})$ is 1-dimensional (Lem. 3.2) and so, as S_D-representations, we have $\mathsf{HWV}_{\mu^*}\left(\bigotimes^D W\right) \simeq [\mu]$, which is irreducible. The linear span of the S_D-orbit of $\langle\widehat{\mu}|$ is an S_D-representation. It is contained in the irreducible S_D-representation $\mathsf{HWV}_{\mu^*}\left(\bigotimes^D W\right)$ and hence equals it. □

4.3.1 The Tensor Scenario

In analogy to the above discussion, we now treat the tensor scenario. For a partition triple $\lambda \vdash_n^* d$ we define

$$\langle\widehat{\lambda}| := \mathsf{reorder}_{3,n}\left(\langle\widehat{\lambda^{(1)}}| \otimes \langle\widehat{\lambda^{(2)}}| \otimes \langle\widehat{\lambda^{(3)}}|\right), \tag{4.5}$$

where the linear isomorphism $\mathsf{reorder}_{a,b}\colon \bigotimes^a\bigotimes^b\mathbb{C}^n \to \bigotimes^b\bigotimes^a\mathbb{C}^n$ is defined by $\bigotimes_{i=1}^a(\bigotimes_{j=1}^b v_{ij}) \mapsto \bigotimes_{j=1}^b(\bigotimes_{i=1}^a v_{ij})$ for $a, b \in \mathbb{N}$. One can check that $\langle\widehat{\lambda}|$ is a HWV of weight λ^*. In analogy to Claim 4.2, we obtain the following two results.

Claim 4.3. $\mathsf{HWV}_{\lambda^*}(\bigotimes^d\bigotimes^3(\mathbb{C}^n)^*)$ *is generated by $\langle\widehat{\lambda}|\pi$ with $\pi \in \mathsf{S}_d^3$.*

Claim 4.4. *The tensors $\langle\widehat{\lambda}|\pi\mathcal{P}_d$ with $\pi \in \mathsf{S}_d^3$ generate* $\mathsf{HWV}_{\lambda^*}(\mathsf{Sym}^d\bigotimes^3(\mathbb{C}^n)^*)$.

Remark 4.5. The tuple (λ, π) succinctly encodes a specific HWV in $\mathsf{HWV}_{\lambda^*}(\bigotimes^d\bigotimes^3(\mathbb{C}^n)^*)$.

5. OBSTRUCTION DESIGNS

In this section we introduce a combinatorial concept, the *obstruction designs*, which enables us to construct HWVs in $\mathsf{Sym}^d\bigotimes^3(\mathbb{C}^n)^*$. Our main tool is Claim 4.4.

5.1 Set Partitions

A *set partition* Λ of a finite set S is a set of subsets of S such that for all $s \in S$ there exists exactly one $e_s \in \Lambda$ with $s \in e_s$. We call e_s the *hyperedge corresponding to* s. The *type* of Λ is defined as the partition $\lambda \vdash |S|$ obtained from sorting the multiset $\{|e| : e \in \Lambda\}$. We write $V(\Lambda) := S$ and put $d := |S|$.

Definition 5.1. An *ordered set partition* Λ of a set S of type λ is a set partition of S of type λ endowed with (1) linear orderings on each hyperedge $e \in \Lambda$ and (2) for each length $1 \le \ell \le \ell(\lambda)$ a linear ordering on the set $\{e \in \Lambda : |e| = \ell\}$ of hyperedges with the cardinality ℓ.

Each Λ induces a numbering $\pi_\Lambda\colon S \to \{1, \ldots, d\}$ of S by sorting the hyperedges according to their size in decreasing order, resolving ties using property (2), and using the ordering of the hyperedges given by property (1), cf. Fig. 1.

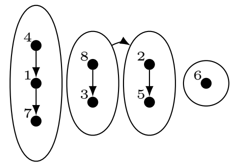

$$= \Lambda \leftrightsquigarrow \pi_\Lambda = \begin{pmatrix} 4 & 1 & 7 & 8 & 3 & 2 & 5 & 6 \\ 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 \end{pmatrix}$$

Figure 1: The bijection $\Lambda \mapsto \pi_\Lambda$ between ordered set partitions of type $\lambda = (4,3,1)$ and S_8 (note ${}^t\lambda = (3,2,2,1)$). The orderings in the left picture are shown with arrows pointing from the smaller element to the bigger one.

It is obvious that $\Lambda \mapsto \pi_\Lambda$ is a bijection between the set of ordered set partitions of S of type λ and the set of bijections $S \to \{1, \ldots, d\}$. In the following we assume that $S = \{1, 2 \ldots, d\}$. Therefore, we can interpret π_Λ as a permutation of S, i.e., $\pi_\Lambda \in S_d$. We denote by Λ^λ the ordered set partition on S corresponding to $\pi_\Lambda = \mathrm{id}$ and by Λ_i^λ its ith hyperedge.

A finite sequence $\zeta = (\zeta_1, \ldots, \zeta_d)$ of vectors ζ_i in $\mathbb{C}^n$ shall be a called *list*. The group S_d acts naturally on the set of lists by permuting the positions as follows:

$$\pi(\zeta_1, \zeta_2, \ldots, \zeta_d) := (\zeta_{\pi^{-1}(1)}, \zeta_{\pi^{-1}(2)}, \ldots, \zeta_{\pi^{-1}(d)}).$$

Given a list $\zeta \in (\mathbb{C}^n)^d$ and a hyperedge e of Λ we define the *restriction* $\zeta_{|_e}$ of ζ to e by

$$(\zeta_1, \zeta_2, \ldots, \zeta_d)_{|_e} := (\zeta_{e^1}, \ldots, \zeta_{e^\ell}),$$

where $e^1, \ldots, e^\ell$ are the elements of e ordered according to the internal ordering of e. If e denotes the ith hyperedge of Λ, then $\zeta_{|_e} = \pi_\Lambda(\zeta)_{|_{\Lambda_i^\lambda}}$.

5.1.1 Highest Weight Vectors

For each ordered set partition Λ of type $\lambda \vdash_n d$ we have $\pi_\Lambda \in \mathsf{S}_d$ and hence obtain a nonzero HWV

$$f_\Lambda := \langle\widehat{\lambda}|\pi_\Lambda \in \bigotimes^d(\mathbb{C}^n)^*$$

of weight λ^*, provided $n \ge \ell(\lambda)$.

Consider the Young subgroup $Y_\lambda^{\text{inner}} := \mathsf{S}(\Lambda_1^\lambda) \times \cdots \times \mathsf{S}(\Lambda_{\lambda_1}^\lambda)$. Let Y_λ^{outer} denote the group that interchanges hyperedges of the same length in Λ^λ while preserving the order in each hyperedge. Further, let Y_λ denote the subgroup of S_d generated by Y_λ^{inner} and Y_λ^{outer}. Then $\langle\widehat{\lambda}|\tau = \pm\langle\widehat{\lambda}|$ for $\tau \in Y_\lambda$. Moreover, the permutations in the coset $Y_\lambda\pi_\Lambda$ correspond to the ordered set partitions that result from Λ be changing the orderings (cf. Def. 5.1). This shows that for a fixed partition $\lambda \vdash d$ there is a bijection between the set of right cosets of $Y_\lambda \subseteq \mathsf{S}_d$ and the set of set partitions of type λ. Hence a set partition Λ of type $\lambda \vdash_n d$ uniquely determines a HWV f_Λ of weight λ^*

$$f_\Lambda := \pm\langle\widehat{\lambda}|\pi_\Lambda \in \bigotimes^d(\mathbb{C}^n)^*$$

up to a sign, where π_Λ is the permutation corresponding to some ordering of Λ.

One can prove that the *projective stabilizer* $Y_\lambda := \{\tau \in \mathsf{S}_d : \langle\widehat{\lambda}|\tau = \pm\langle\widehat{\lambda}|\}$ is generated by Y_λ^{inner} and Y_λ^{outer}.

5.1.2 Contraction

For a list $\zeta = (\zeta_1, \dots, \zeta_d)$ we write $|\zeta\rangle := |\zeta_1 \otimes \cdots \otimes \zeta_d\rangle$. We are going to analyze how the scalar product $\langle\widehat{\lambda}|\pi|\zeta\rangle$ can be interpreted combinatorially using set partitions, for $\pi \in \mathsf{S}_d$.

Setting $\vartheta := \pi\zeta$ we obtain from (4.4)

$$\langle\widehat{\lambda}|\pi|\zeta\rangle = \langle\widehat{\lambda}|\vartheta\rangle = \textstyle\prod_{i=1}^{\lambda_1} \langle\widehat{{}^t\lambda_i}|\vartheta_{|\Lambda_i^\lambda}\rangle,$$

where we recall that Λ_i^λ is the ith hyperedge of Λ^λ. Note that $\langle\widehat{\ell}|\vartheta\rangle$ is just the determinant of the $\ell \times \ell$-matrix $\big(\langle i|\vartheta_j\rangle\big)_{i,j}$, for a list $\vartheta = (\vartheta_1, \dots, \vartheta_\ell)$ with $\ell \le n$, see (4.3).

Fix an ordered set partition Λ. Given a hyperedge $e \in \Lambda$ and a hyperedge labeling $\zeta^e\colon e \to \mathbb{C}^n$, we write $|\zeta^e\rangle := |\zeta^{e^1}\rangle \otimes \cdots \otimes |\zeta^{e^\ell}\rangle$, where $e^1, \dots, e^\ell$ are the elements of e in their order. We define the *evaluation* $\mathsf{eval}_e(\zeta^e) := \langle\widehat{\ell}|\zeta^e\rangle \in \mathbb{C}$. Note that the evaluation $\mathsf{eval}_e(\zeta^e)$ is, up to sign, invariant under changing the linear order of e. For a labeling $\zeta\colon V(\Lambda) \to \mathbb{C}^n$ we define the *evaluation of the ordered set partition* Λ *at the labeling* ζ:

$$\mathsf{eval}_\Lambda(\zeta) := \textstyle\prod_{e\in\Lambda} \mathsf{eval}_e(\zeta_{|_e}).$$

We therefore conclude that $\mathsf{eval}_\Lambda(\zeta) = \langle\widehat{\lambda}|\pi_\Lambda|\zeta\rangle$.

5.2 Obstruction Designs

An *ordered set partition triple* $\mathcal{H}$ consists of a vertex set $V(\mathcal{H}) = \{1, \dots, d\}$ and three ordered set partitions $E^{(k)} := E^{(k)}(\mathcal{H})$, $k \in \{1,2,3\}$, of $V(\mathcal{H})$. The elements of each $E^{(k)}$ are called *hyperedges*. The ordered set partition triple $\mathcal{H}$ is said to have *type* λ, where $\lambda \vdash^{\!*} d$ is a partition triple, if the set partition $E^{(k)}$ has type $\lambda^{(k)}$ for all k. Via our explicit bijections between S_d and the set of ordered set partitions of a fixed type, we get an explicit bijection between S_d^3 and the set of ordered set partition triples of type λ. The permutation triple corresponding to an ordered set partition $\mathcal{H}$ is denoted by $\pi_\mathcal{H}$.

For partition triples $\lambda \vdash^{\!*} d$ define $Y_\lambda := Y_{\lambda^{(1)}} \times Y_{\lambda^{(2)}} \times Y_{\lambda^{(3)}}$, where $Y_{\lambda^{(k)}}$ is the projective stabilizer defined in Sec. 5.1. We can again forget about the orderings and arrive at the following definition.

Definition 5.2. A *set partition triple* $\mathcal{H}$ consists of a vertex set $V(\mathcal{H}) = \{1, \dots, d\}$ and three set partitions $E^{(1)}, E^{(2)}, E^{(3)}$ of $V(\mathcal{H})$.

The above discussion implies , analogously to Subsec. 5.1:

Claim 5.3. *For a fixed partition triple $\lambda \vdash^{\!*} d$ there is a bijection between the right cosets of $Y_\lambda \subseteq \mathsf{S}_d^3$ and the set of set partition triples of type λ.*

So a set partition triple $\mathcal{H}$ defines (up to sign) the HWV

$$\pm\langle\widehat{\lambda}|\pi_\mathcal{H} \in \textstyle\bigotimes^d \bigotimes^3 (\mathbb{C}^n)^*$$

of weight λ^*, where $\lambda \vdash^{\!*}_{n} d$ denotes the type of $\mathcal{H}$.

One can show that Y_λ is the *projective stabilizer* of $\langle\widehat{\lambda}|$, i.e., $Y_\lambda = \{\tau \in \mathsf{S}_d^3 : \langle\widehat{\lambda}|\tau = \pm\langle\widehat{\lambda}|\}$.

5.2.1 Triple Contraction

A finite sequence of vectors in $(\mathbb{C}^n)^3$ shall be called a *triple list*. Given a triple list ζ containing d triples, we write

$$\zeta = \begin{pmatrix} \zeta_1^{(1)}, \dots, \zeta_d^{(1)} \\ \zeta_1^{(2)}, \dots, \zeta_d^{(2)} \\ \zeta_1^{(3)}, \dots, \zeta_d^{(3)} \end{pmatrix}, \quad \begin{aligned} \zeta^{(k)} &:= (\zeta_1^{(k)}, \dots, \zeta_d^{(k)}), \\ \zeta_i &:= (\zeta_i^{(1)}, \zeta_i^{(2)}, \zeta_i^{(3)}). \end{aligned}$$

Moreover, we write $|\zeta\rangle := \mathsf{reorder}_{3,d}(|\zeta^{(1)}\rangle \otimes |\zeta^{(2)}\rangle \otimes |\zeta^{(3)}\rangle) \in \bigotimes^d \bigotimes^3 \mathbb{C}^n$. For an ordered subset $e \subseteq V(\mathcal{H})$ of vertices we identify triple lists $(\zeta_1, \dots, \zeta_{|e|})$ with labelings on e whose codomain is $(\mathbb{C}^n)^3$.

Given a labeling $\zeta\colon V(\mathcal{H}) \to (\mathbb{C}^n)^3$ of an ordered set partition triple $\mathcal{H}$, we define the evaluation of $\mathcal{H}$ at ζ as follows:

$$\mathsf{eval}_\mathcal{H}(\zeta) := \textstyle\prod_{k=1}^3 \mathsf{eval}_{E^{(k)}(\mathcal{H})}(\zeta^{(k)}).$$

Then we have, if λ denotes the type of $\mathcal{H}$,

$$\mathsf{eval}_\mathcal{H}(\zeta) = \langle\widehat{\lambda}|\pi_\mathcal{H}|\zeta\rangle. \tag{5.1}$$

5.2.2 Symmetrization

Since $\tau\mathcal{P}_d = \mathcal{P}_d$ for all $\tau \in \mathsf{S}_d$, we get $\langle\widehat{\lambda}|\pi\mathcal{P}_d = \langle\widehat{\lambda}|\pi\tau\mathcal{P}_d$ for all $\tau \in \mathsf{S}_d$. Hence the polynomial

$$f_\mathcal{H} := \langle\widehat{\lambda}|\pi_\mathcal{H}\mathcal{P}_d \in \mathsf{Sym}^d \textstyle\bigotimes^3 (\mathbb{C}^n)^* \tag{5.2}$$

(cf. Lem. 4.1) assigned to a set partition triple $\mathcal{H}$ is independent of the numbering of its vertices. This motivates the following crucial definition. (For a similar graph construction see [14, Sec. 6.9].)

Definition 5.4. An *obstruction design* is defined to be an equivalence class of set partition triples under renumbering of the vertices, satisfying $|e_1 \cap e_2 \cap e_3| \le 1$ for all hyperedge triples $(e_1, e_2, e_3) \in E^{(1)} \times E^{(2)} \times E^{(3)}$.

The restriction on $|e_1 \cap e_2 \cap e_3|$ in Def. 5.4 is motivated by the following lemma.

Lemma 5.5. *If an ordered set partition $\mathcal{H}$ contains two vertices y and y' which lie in the same three hyperedges, then $f_\mathcal{H}$ is the zero polynomial.*

Each obstruction design $\mathcal{H}$ describes the polynomial $f_\mathcal{H}$ of degree d, up to sign. Since we do not care about the sign, we abuse notation in the following way: For every $\mathcal{H}$ we implicitly fix an ordered set partition triple $\mathcal{H}'$ such that $\mathcal{H}$ is obtained from $\mathcal{H}'$ by forgetting about orderings and vertex numbers. Then we define $\mathsf{eval}_\mathcal{H}(\zeta) := \mathsf{eval}_{\mathcal{H}'}(\zeta)$.

Corollary 5.6. *Let $\mathcal{H}$ be an ordered set partition triple with d vertices and $\xi\colon V(\mathcal{H}) \to (\mathbb{C}^n)^3$ be a labeling. We have*

$$\langle\widehat{\lambda}|\pi_\mathcal{H}\mathcal{P}_d|\xi\rangle = \frac{1}{d!} \sum_{\zeta \in \mathsf{S}_d \xi} \mathsf{eval}_\mathcal{H}(\zeta)$$

Proof. Follows from (5.1) and the definition of $\mathcal{P}_d$ in (4.2). □

Theorem 5.7. *For a partition triple λ we have*

$$\mathsf{HWV}_{\lambda^*}(\mathsf{Sym}^d \textstyle\bigotimes^3 (\mathbb{C}^*)^n) = \operatorname{span}\{f_\mathcal{H} : \mathcal{H} \textit{ obs. des. of type } \lambda\}.$$

In particular, $k(\lambda) = \dim \operatorname{span}\{f_\mathcal{H} : \mathcal{H} \textit{ obs. des. of type } \lambda\}$ *and* $k(\lambda) \le |\{\mathcal{H} : \mathcal{H} \textit{ obs. des. of type } \lambda\}|$.

Proof. According to Claim 4.4, for $\lambda \vdash^{\!*}_{n} d$, we have $\mathsf{HWV}_{\lambda^*}(\mathsf{Sym}^d \bigotimes^3 (\mathbb{C}^*)^n) = \operatorname{span}\{\langle\widehat{\lambda}|\pi_\mathcal{H}\mathcal{P}_d : \pi \in \mathsf{S}_d^3\}$. But since obstruction designs $\mathcal{H}$ determine $f_\mathcal{H} = \langle\widehat{\lambda}|\pi_\mathcal{H}\mathcal{P}_d$ up to a sign, the first assertion follows. Now we use the definition of the Kronecker coefficients. □

5.2.3 Kronecker Coeff. and Discrete Tomography

We identify a finite subset $r \subseteq \mathbb{N} \times \mathbb{N} \times \mathbb{N}$ with its characteristic function $r\colon \mathbb{N} \times \mathbb{N} \times \mathbb{N} \to \{0, 1\}$ and define its *first marginal distribution* $p^{(1)}(r)$ as $p^{(1)}(r) := \big(\sum_{j,l \in \mathbb{N}} r(i, j, l)\big)_{i \in \mathbb{N}}$. Analogously, we define $p^{(2)}$ and $p^{(3)}$. For a given partition triple λ let

$$\mathscr{R}_\lambda := \{r \subseteq \mathbb{N} \times \mathbb{N} \times \mathbb{N} \mid p^{(k)}(r) = {}^t\lambda^{(k)} \text{ for all } 1 \le k \le 3\}$$

denote the set of subsets $r \subseteq \mathbb{N} \times \mathbb{N} \times \mathbb{N}$ with marginal distributions prescribed by the transpose of $\lambda^{(k)}$. Note that all triples in $r \in \mathscr{R}_\lambda$ are contained in a box of side lengths $b_1 \times b_2 \times b_3$, where $b_k := \lambda_1^{(k)}$.

We define a surjective map φ from $\mathscr{R}_\lambda$ to the set of obstruction designs of type λ as follows: To $r \in \mathscr{R}_\lambda$ we assign $\mathcal{H}$ such that $V(\mathcal{H}) := r$ and $E^{(1)} := \{(a, j, l) : 1 \le j \le \lambda_1^{(2)},\ 1 \le l \le \lambda_1^{(3)}\}$, with analogous definitions for $E^{(2)}, E^{(3)}$. It is easily verified that $\mathcal{H}$ is indeed an obstruction design of type λ.

The group $\mathsf{S}_{b_1} \times \mathsf{S}_{b_2} \times \mathsf{S}_{b_3}$ acts on $\mathbb{N}^{b_1} \times \mathbb{N}^{b_2} \times \mathbb{N}^{b_3}$ via permutation of positions. The stabilizer of $({}^t\lambda^{(1)}, {}^t\lambda^{(2)}, {}^t\lambda^{(3)})$ in $\mathsf{S}_{b_1} \times \mathsf{S}_{b_2} \times \mathsf{S}_{b_3}$ acts on $\mathscr{R}_\lambda$ in the natural way. It is readily checked that $\varphi(r_1) = \varphi(r_2)$ iff r_1 and r_2 lie in the same orbit w.r.t. this operation. Hence the number of orbits $|\mathscr{R}_\lambda/{\sim}|$ equals the number of obstruction designs of type λ. We conclude from Thm. 5.7 the following result (which is related, but different from [22]).

Corollary 5.8. $k(\lambda) \le |\mathscr{R}_\lambda/{\sim}| \le |\mathscr{R}_\lambda|$.

Consider the following decision problem: Given a partition triple $\lambda \vDash^{*} d$ encoded in unary. [1] proved that the problem of deciding $|\mathscr{R}_{{}^t\lambda}| > 0$ is NP-complete. Since $|\mathscr{R}_\lambda/{\sim}| > 0$ iff $|\mathscr{R}_\lambda| > 0$, we obtain:

Proposition 5.9. *Given a partition triple λ encoded in unary. Then it is NP-complete to decide whether there exists an obstruction design of type λ.*

5.2.4 Polynomial Evaluation

The polynomial $f_\mathcal{H}$ corresponding to an obstruction design $\mathcal{H}$, cf. (5.2), can be evaluated at a point $|w\rangle = \sum_{i=1}^{r} |w_i^{(1)}\rangle \otimes |w_i^{(2)}\rangle \otimes |w_i^{(3)}\rangle \in \bigotimes^3 \mathbb{C}^n$ as follows:

$$\begin{aligned} f_\mathcal{H}(w) &\overset{\text{Lem. 4.1}}{=} \langle \widehat{\lambda} | \pi_\mathcal{H} \mathcal{P}_d | w^{\otimes d} \rangle = \langle \widehat{\lambda} | \pi_\mathcal{H} | w^{\otimes d} \rangle \\ &= \textstyle\sum_{J \in \{1,\dots,r\}^d} \langle \widehat{\lambda} | \pi_\mathcal{H} | w_{J_1} w_{J_2} \cdots w_{J_d} \rangle \\ &\overset{(5.1)}{=} \textstyle\sum_{J \in \{1,\dots,r\}^d} \mathsf{eval}_\mathcal{H}(w_{J_1}, w_{J_2}, \dots, w_{J_d}). \end{aligned} \tag{5.3}$$

We can interpret $\zeta := (w_{J_1}, w_{J_2}, \dots, w_{J_d})$ as a vertex labeling. Hence the above sum is over all vertex labelings $\zeta\colon V(\mathcal{H}) \to \{w_1, \dots, w_r\}$, where we interpret the codomain as a multiset.

Example 5.10. Let $f \in \mathsf{Sym}^{nl} \bigotimes^3 (\mathbb{C}^n)^*$ be a HWV of weight λ^* such that $\lambda^{(1)} = \lambda^{(2)} = \lambda^{(3)} = n \times l$ with l odd and $n > 1$. Then f vanishes on $\mathsf{GL}_n^3 \mathcal{E}_n$.

The following fundamental questions arise when studying obstruction designs.

Questions 5.11. *(1) Given an obstruction design $\mathcal{H}$ of type $\lambda \vdash^{*}_{n} d$ and a tensor $v \in \bigotimes^3 \mathbb{Z}^n$. What is the complexity of computing the evaluation $f_\mathcal{H}(v)$? Is this problem $\#\mathsf{P}$-hard under Turing reductions?*

(2) Given an obstruction design $\mathcal{H}$ of type $\lambda \vdash^{*}_{n} d$. What is the complexity of deciding whether $f_\mathcal{H} = 0$?

(3) For a given partition triple $\lambda \vdash^{*}_{n} d$, explicitly describe a maximal linear independent subset of the set of obstruction designs of type λ!

An answer to Question 5.11(3) would result in an explicit basis of $([\lambda^{(1)}] \otimes [\lambda^{(2)}] \otimes [\lambda^{(3)}])^{\mathsf{S}_d}$ and solve one of the most fundamental open questions in the representation theory of the symmetric groups.

6. LOWER BOUNDS FOR MATRIX MULTIPLICATION

For $\kappa \in \mathbb{N}$, we define the partition triple λ consisting of the three hooks $(2\kappa + 1) \ulcorner (\kappa + 1)$ of size $d := 3\kappa + 1$ and length $2\kappa + 1$.

We construct an obstruction design $\mathcal{H}$ of type λ as follows (see Fig. 2). The vertex set $V(\mathcal{H})$ is partioned into disjoint sets $V^{(1)} \mathbin{\dot\cup} V^{(2)} \mathbin{\dot\cup} V^{(3)} \mathbin{\dot\cup} \{y^0\}$, where $|V^{(k)}| = \kappa$ for all k. Each $E^{(k)}$ consists of one hyperedge $e^{(k)} := V^{(k+1)} \cup V^{(k+2)} \cup \{y^0\}$ of size $2\kappa + 1$ (addition mod 3 in the exponent) and κ singletons.

One can prove that $\mathcal{H}$ is the only obstruction design of type λ, hence $k(\lambda) \le 1$. (Note that $k(\lambda) = 1$ by [18, Thm. 2.1], so we must have $f_\mathcal{H} \neq 0$ by Thm. 5.7.)

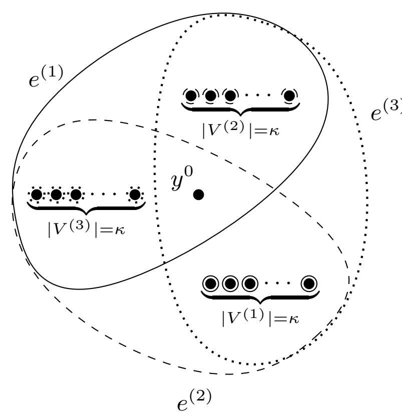

Figure 2: The unique family of obstruction designs corresponding to the hook partition triple λ.

Proposition 6.1. *Fix an odd number $m \in \mathbb{N}$ and set $\kappa := \frac{m^2-1}{2}$ and $d := 3\kappa + 1$.*

(1) All matrix triples $B \in (\mathbb{C}^{3\kappa \times 3\kappa})^3$ satisfy $f_\mathcal{H}(B\mathcal{E}_{3\kappa}) = 0$.

(2) There exists a matrix triple $A \in (\mathbb{C}^{m^2 \times m^2})^3$ such that $f_\mathcal{H}(A\mathcal{M}_m) \neq 0$.

Note that $f_\mathcal{H}$ is a homogeneous polynomial of degree d on $\bigotimes^3 \mathbb{C}^{m^2}$. Prop. 6.1 explicitly exhibits an obstruction family and directly implies the following Thm. 6.2 for odd numbers.

Theorem 6.2.

$$\underline{R}(\mathcal{M}_m) \ge \begin{cases} \frac{3}{2}m^2 - 2 & \text{for } m \text{ even} \\ \frac{3}{2}m^2 - \frac{1}{2} & \text{for } m \text{ odd} \end{cases}.$$

We omit to handle the case where m is even.

6.1 Vanishing on the Unit Tensor Orbit

In this subsection we prove Prop. 6.1(1).

Let $B \in (\mathbb{C}^{3\kappa\times 3\kappa})^3$ be arbitrary. We define the triple list

$$w := \big((B^{(1)}|1\rangle, B^{(2)}|1\rangle, B^{(3)}|1\rangle), \ldots,\\ (B^{(1)}|3\kappa\rangle, B^{(2)}|3\kappa\rangle, B^{(3)}|3\kappa\rangle)\big).$$

According to (5.3) we have

$$f_{\mathcal{H}}(\mathcal{E}_{3\kappa}) = \sum_{J\in\{1,\ldots,3\kappa\}^{3\kappa+1}} \mathsf{eval}_{\mathcal{H}}(Bw_{J_1}, \ldots, Bw_{J_{3\kappa+1}}). \quad (*)$$

The crucial property of $\mathcal{H}$ is that for each pair of vertices $\{y_1, y_2\}$ there exists a hyperedge e of $\mathcal{H}$ containing both y_1 and y_2. By the pidgeon-hole principle, for each labeling $J\colon V(\mathcal{H}) \to \{1, \ldots, 3\kappa\}$ there exists a pair of vertices $\{y_1, y_2\}$ such that $J(y_1) = J(y_2)$. Hence, by the crucial property of $\mathcal{H}$, y_1 and y_2 lie in a common hyperedge e. Therefore, $\mathsf{eval}_e((Bw_{J_1}, \ldots, Bw_{J_{3\kappa+1}})_{|_e}) = 0$, because the latter is the determinant of a matrix with two columns equal. Hence each summand in $(*)$ vanishes, proving Prop. 6.1(1).

6.2 Evaluation at the Matrix Mult. Tensor

In this subsection we prove Prop. 6.1(2). For notational convenience, we define the triples (omitting parentheses)

$$t_{ijl} := \big(|ij\rangle, |jl\rangle, |li\rangle\big) \in (\mathbb{C}^{m\times m})^3 \tag{6.1}$$

and the triple list w of length m^3 obtained by concatenating all t_{ijl} for $1 \le i, j, l \le m$ in any order. We put $\mathscr{T} := \{t_{ijl} \mid 1 \le i, j, l \le m\}$. Recall from (2.1):

$$\mathcal{M}_m = \sum_{i,j,l=1}^{m} t_{ijl}^{(1)} \otimes t_{ijl}^{(2)} \otimes t_{ijl}^{(3)}.$$

The strategy is to construct an $m^2\times m^2$ matrix triple A with affine linear matrix entries in indeterminates $X_1, \ldots, X_N$. According to (5.3) we have

$$f_{\mathcal{H}}(A\mathcal{M}_m) = \textstyle\sum_{J\in\{1,\ldots,m^3\}^d} \mathsf{eval}_{\mathcal{H}}(Aw_{J_1}, \ldots, Aw_{J_d}). \quad (\dagger)$$

We shall exhibit a monomial $\mathcal{X}$ in the X_i whose coefficient in $f_{\mathcal{H}}(A\mathcal{M}_m)$ is nonzero. Hence $f(X) := f_{\mathcal{H}}(A\mathcal{M}_m)$ is not the zero polynomial in the X_i. There is a substitution of the X_i with suitable values $\alpha_1, \ldots, \alpha_N \in \mathbb{C}$ such that $f(\alpha) \ne 0$. Making this substitution in A yields the desired matrix triple over $\mathbb{C}$.

6.2.1 *Invariance in each* $V^{(k)}$

We use the short notation $\mathsf{eval}_e(\zeta) := \mathsf{eval}_e(\zeta^{(k)}_{|_e})$ for a hyperedge $e \in E^{(k)}$ and a triple labeling ζ.

CLAIM 6.3. *Let $\sigma\colon V(\mathcal{H}) \to V(\mathcal{H})$ be a bijection satisfying $\sigma(V^{(k)}) = V^{(k)}$ for all $k \in \{1, 2, 3\}$. For every triple labeling $\zeta\colon V(\mathcal{H}) \to (\mathbb{C}^{m^2})^3$ we have $\mathsf{eval}_{\mathcal{H}}(\zeta) = \mathsf{eval}_{\mathcal{H}}(\zeta\circ\sigma)$.*

PROOF. It suffices to show the claim for a transposition σ exchanging two elements of $V^{(1)}$, because the situation for $V^{(2)}$ and $V^{(3)}$ is completely symmetric. We have $\prod_{e\in E^{(1)}} \mathsf{eval}_e(\zeta) = \prod_{e\in E^{(1)}} \mathsf{eval}_e(\zeta\circ\sigma)$, because, up to reordering, both products have the same factors. For $k \in \{2, 3\}$ we have $\mathsf{eval}_e(\zeta) = \mathsf{eval}_e(\zeta\circ\sigma)$ for every singleton hyperedge $e \in E^{(k)}$ and $\mathsf{eval}_{e^{(k)}}(\zeta) = -\mathsf{eval}_{e^{(k)}}(\zeta\circ\sigma)$. Therefore $\prod_{e\in E^{(k)}} \mathsf{eval}_e(\zeta) = -\prod_{e\in E^{(k)}} \mathsf{eval}_e(\zeta\circ\sigma)$. As a result we get $\mathsf{eval}_{\mathcal{H}}(\zeta) = (-1)^2 \mathsf{eval}_{\mathcal{H}}(\zeta\circ\sigma)$. $\square$

6.2.2 *Special Structure of the Matrix Triple*

Let $\Gamma := \mathbb{C}[X_i^{(k)} : 1 \le k \le 3,\ 1 \le i \le m]$ denote the polynomial ring in $3m$ variables. Recall that m is odd and $\kappa = \frac{m^2-1}{2}$. We set $\bar{i} := m + 1 - i$ for $1 \le i \le m$, thinking of $i \mapsto \bar{i}$ as a reflection at $a := (m+1)/2$. Note $\bar{a} = a$. We consider the set of pairs $O_m := \{1, \ldots, m\}\times\{1, \ldots, m\}\setminus\{aa\}$ and fix an arbitrary bijection $\varphi\colon O_m \to \{2, \ldots, m^2\}$.

For each $1 \le k \le 3$ we define the matrix $A^{(k)}$ of format $(m\times m)\times m^2$ with the following affine linear entries in $X_i^{(k)}$:

$$A^{(k)}|ij\rangle := \begin{cases} X_a^{(k)}|1\rangle & \text{if } i = j = a \\ |\varphi(i\bar{i})\rangle + X_i^{(k)}|1\rangle & \text{if } i \ne j \text{ and } j = \bar{i}\ . \\ |\varphi(ij)\rangle & \text{if } j \ne \bar{i} \end{cases}$$

Hence $A^{(k)}$ looks as follows:

$$\left(\begin{array}{ccccccc|c} X_a^{(k)} & X_1^{(k)} & \cdots & X_{a-1}^{(k)} & X_{a+1}^{(k)} & \cdots & X_m^{(k)} & \\ & 1 & & & & & & \\ & & \ddots & & & & & 0 \\ & & & 1 & & & & \\ & & & & 1 & & & \\ & & & & & \ddots & & \\ & & & & & & 1 & \\ \hline & & & 0 & & & & \mathrm{id}_{m^2-m} \end{array}\right), \quad (※)$$

where we arranged the rows and columns as follows: The left m columns correspond to the vectors $|i\bar{i}\rangle$, where the leftmost one corresponds to $|aa\rangle$. The top row corresponds to the vector $|1\rangle$ and the following $m-1$ rows correspond to the vectors $|\varphi(i\bar{i})\rangle$. Recall that $f_{\mathcal{H}}(A\mathcal{M}_m)$ is a sum of products of determinants of submatrices of the $A^{(k)}$.

The sum $f_{\mathcal{H}}(A\mathcal{M}_m)$ is an element of Γ and we are interested in its coefficient of the monomial $\mathcal{X}$, where

$$\mathcal{X} := \prod_{k=1}^{3} X_a^{(k)} \prod_{i=1}^{m} \big(X_i^{(k)}\big)^{|i-\bar{i}|}. \tag{6.2}$$

We remark that the degree of $\mathcal{X}$ is $3(1 + \sum_{i=1}^{m} |i - \bar{i}|)$. It is readily checked that $\sum_{i=1}^{m} |i - \bar{i}| = \kappa$.

Fix any numbering of the vertices of $\mathcal{H}$. For $J \in \{1, \ldots, m^3\}^d$ we abuse notation and define the map $J\colon V(\mathcal{H}) \to \mathscr{T}$ via $J(y) := w_{J_y}$. With this notation, $(\dagger)$ becomes $\sum_J \mathsf{eval}_{\mathcal{H}}\big(AJ(1), \ldots, AJ(d)\big)$, or $\sum_J \mathsf{eval}_{\mathcal{H}}(AJ)$ in short notation. We call a triple labeling $J\colon V(\mathcal{H}) \to \mathscr{T}$ *nonzero*, if the coefficient of $\mathcal{X}$ in the polynomial $\mathsf{eval}_{\mathcal{H}}(AJ)$ is nonzero. We will count and classify all nonzero triple labelings J and show that all $\mathsf{eval}_{\mathcal{H}}(AJ)$ contribute the same coefficient with respect to the monomial $\mathcal{X}$. This implies that the coefficient of $\mathcal{X}$ in $f_{\mathcal{H}}(A\mathcal{M}_m)$ is a sum without cancellations and hence is nonzero.

6.2.3 *Separate Analysis of the Three Layers*

We fix a nonzero triple labeling $J\colon V(\mathcal{H}) \to \mathscr{T}$ and write $J = (J^{(1)}, J^{(2)}, J^{(3)})$. Recall that the hyperedge $e^{(k)}$ has size $2\kappa + 1 = m^2$. Since J is nonzero, $J^{(k)}$ is injective on hyperedges and therefore $|\{J^{(k)}(y) : y \in e^{(k)}\}| = m^2$. Hence $J^{(k)}$ is bijective on $e^{(k)}$, as $|J^{(k)}(V(\mathcal{H}))| \le m^2$.

CLAIM 6.4. *For all $y \in V^{(k)}$ we have $J^{(k)}(y) = |i\bar{i}\rangle$ for some $1 \le i \le m$.*

PROOF. Since $\{y\} \in E^{(k)}$ and J is nonzero, we have $\langle 1|A^{(k)}|J^{(k)}(y)\rangle \neq 0$. From the definition of A it follows that $J^{(k)}(y) = |ij\rangle$ and the third case $j \neq \bar{i}$ is excluded. Hence $j = \bar{i}$. $\square$

CLAIM 6.5. *We have* $J(y^0) = (|aa\rangle, |aa\rangle, |aa\rangle)$.

PROOF. For the following argument it is important to keep the structure of the matrix $A^{(k)}$ in mind, cf. (※). Recall that $f_{\mathcal{H}}(A\mathcal{M}_m)$ is a sum of products of certain subdeterminants of $A^{(k)}$ that are determined by the hyperedges in $E^{(k)}(\mathcal{H})$. The coefficient of $\mathcal{X}$ in $\mathsf{eval}_{\mathcal{H}}(AJ(1), \ldots, AJ(d))$ is nonzero as J is nonzero. Fix k. Since the degree of $X_a^{(k)}$ in $\mathcal{X}$ is one, there is exactly one vertex $y_k \in V(\mathcal{H})$ with $J^{(k)}(y) = |aa\rangle$. But we know that $J^{(k)}$ bijective on $e^{(k)}$, so $y_k \in e^{(k)}$.

It is now sufficient to show that $y_1 = y_2 = y_3$ (since $e^{(1)} \cap e^{(2)} \cap e^{(3)} = \{y^0\}$).

The structure of the matrix multiplication tensor implies that $J(y_1) = (|aa\rangle, |ai\rangle, |ia\rangle)$ for some $1 \leq i \leq m$.

In the case $a = i$, by definition of y_2 and y_3 and uniqueness, we have $y_1 = y_2 = y_3$ and we are done.

So consider the case where $a \neq i$. If $y_1 \neq y^0$ we may assume w.l.o.g. $y_1 \in V^{(3)}$. Using Claim 6.4 we conclude that $J^{(3)}(y_1) = |i\bar{i}\rangle$ for some $1 \leq i \leq m$. Hence $\bar{i} = a$ contradicting $i \neq a$. So we must have $y_1 = y^0$.

Similarly, we show that $y_2 = y_3 = y^0$ and the assertion follows. $\square$

CLAIM 6.6. *We have* $J^{(k)}(V^{(k)}) = \{|i\bar{i}\rangle \mid 1 \leq i \leq m\} \setminus \{|aa\rangle\}$, *where the preimage of each* $|i\bar{i}\rangle$ *under* $J^{(k)}$ *has size* $|i - \bar{i}|$.

PROOF. According to Claim 6.5 we have $J(y^0) = (|aa\rangle, |aa\rangle, |aa\rangle)$. Since $A^{(k)}|aa\rangle$ is a multiple of $|1\rangle$, $\mathsf{eval}_{e^{(k)}}(J)$ is a multiple of $X_a^{(k)}$, cf. (※). Moreover, for $i \neq a$, the variable $X_i^{(k)}$ does not appear in the expansion of $\mathsf{eval}_{e^{(k)}}(J^{(k)})$. Since there are $\kappa = \sum_{i=1}^m |i - \bar{i}|$ many contributions of a factor $X_i^{(k)}$ in the monomial $\mathcal{X}$, these factors must be contributed at vertices in $V^{(k)}$. Moreover $|V^{(k)}| = \kappa$, so the only possibility is that all $y \in V^{(k)}$ satisfy $J^{(k)}(y) = |i\bar{i}\rangle$ for some $1 \leq i \leq m$, $i \neq a$. The specific requirement for the number of factors $X_i^{(k)}$ which are encoded in $\mathcal{X}$ in (6.2) finishes the proof. $\square$

6.2.4 Coupling the Analysis of the Three Layers

It will be convenient to identify the sets $J^{(k)}(V^{(k')})$ with their corresponding subsets of O_m.

Consider the bijective map $\tau\colon O_m \to O_m$, $\tau(ij) = (j\bar{i})$, which corresponds to the rotation by 90°. Clearly, $\tau^4 = \mathrm{id}$. The map τ induces a map $\wp(O_m) \to \wp(O_m)$ on the powerset, which we also denote by τ.

Taking the complement defines the involution $\iota\colon \wp(O_m) \to \wp(O_m)$, $S \mapsto O_m \setminus S$. Clearly, we have $\tau \circ \iota = \iota \circ \tau$. We will only be interested in subsets $S \subseteq O_m$ with exactly $|O_m|/2 = \kappa$ many elements and their images under τ and ι. The subsets $S \subseteq O_m$ that satisfy $\iota(S) = \tau(S)$ will be of special interest. Geometrically, these are the sets that get inverted when rotating by 90°.

In Claim 6.6 we analyzed the labels $J^{(k)}(V^k)$. In the next claim we turn to $J^{(k)}(V^{k'})$, where $k \neq k'$.

CLAIM 6.7. *Every nonzero triple labeling J is completely determined by the image $J^{(1)}(V^{(3)})$ (up to permutations in the $V^{(k)}$, see Claim 6.3) as follows.*

- $J^{(2)}(V^{(3)}) = \tau(J^{(1)}(V^{(3)}))$,
- $J^{(2)}(V^{(1)}) = \iota(J^{(2)}(V^{(3)}))$,
- $J^{(3)}(V^{(1)}) = \tau(J^{(2)}(V^{(1)}))$,
- $J^{(3)}(V^{(2)}) = \iota(J^{(3)}(V^{(1)}))$,
- $J^{(1)}(V^{(2)}) = \tau(J^{(3)}(V^{(2)}))$.

Moreover, $\tau(J^{(1)}(V^{(3)})) = \iota(J^{(1)}(V^{(3)}))$.

PROOF. According to Claim 6.6, each vertex $y \in V^{(3)}$ satisfies

$$J(y) = \big(|ij\rangle, |\tau(ij)\rangle, |\bar{i}i\rangle\big)$$

for some $1 \leq i, j \leq m$, $i \neq a$. In particular,

$$\tau(J^{(1)}(V^{(3)})) = J^{(2)}(V^{(3)}).$$

Recall that $J^{(2)}$ is bijective on $e^{(2)}$. Using $e^{(2)} = V^{(1)} \dot\cup V^{(3)} \dot\cup \{y^0\}$ we see that

$$J^{(2)}(V^{(1)}) = O_m \setminus J^{(2)}(V^{(3)}) = \iota(J^{(2)}(V^{(3)})).$$

For the same reason, we can deduce $J^{(3)}(V^{(1)}) = \tau(J^{(2)}(V^{(1)}))$ and $J^{(3)}(V^{(2)}) = \iota(J^{(3)}(V^{(1)}))$. And applying these arguments one more time we get $J^{(1)}(V^{(2)}) = \tau(J^{(3)}(V^{(2)}))$ and $J^{(1)}(V^{(3)}) = \tau(J^{(1)}(V^{(2)}))$. Summarizing (recall $\tau \circ \iota = \iota \circ \tau$) we have

$$J^{(1)}(V^{(3)}) = \tau^3\iota^3(J^{(1)}(V^{(3)})) = \tau^{-1}\iota(J^{(1)}(V^{(3)})),$$

which is equivalent to $\tau(J^{(1)}(V^{(3)})) = \iota(J^{(1)}(V^{(3)}))$. $\square$

Definition 6.8. A subset $S \subseteq O_m$ is called *valid*, if

(1) $|S| = \frac{m^2-1}{2} = \kappa$,

(2) $\tau(S) = \iota(S)$,

(3) $|p^{-1}(i)| = |i - \bar{i}|$ for all $1 \leq i \leq m$

where $p\colon S \to \{1, \ldots, m\}$ is the projection to the first component.

PROPOSITION 6.9. *$J^{(1)}(V^{(3)})$ is a valid set for all nonzero triple labelings J. On the other hand, for every valid set S there exists exactly one nonzero triple labeling J with $J^{(1)}(V^{(3)}) = S$, up to permutations in the $V^{(k)}$.*

PROOF. For the first statement, property (2) of Def. 6.8 follows from Claim 6.7 and property (3) of Def. 6.8 follows from Claim 6.6. The second statement can be readily checked with Claim 6.5 and Claim 6.7. $\square$

Figure 3 gives an example for the case $m = 9$. Vertices that appear in all valid sets are drawn with a solid border. Vertices that appear in no valid set are drawn with a dotted border. Vertices that appear in half of all valid sets are drawn with a dashed border. These contain a vertex label x_i or $\overline{x_i}$. Each valid set corresponds to a choice vector $x \in \{\text{true}, \text{false}\}^4$ determining whether the x_i or the $\overline{x_i}$ are contained in S. This results in $2^4 = 16$ valid sets $S \subseteq O_m$.

The next claim classifies all valid sets.

LEMMA 6.10. *A set $S \subseteq O_m$ is valid iff the following conditions are all satisfied (see Figure 3 for an illustration):*

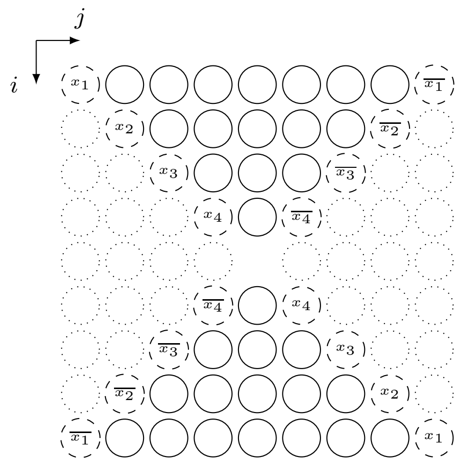

Figure 3: The case $n = 9$.

(1) $\Big\{(ij) \mid (i < j \text{ and } i < \bar{j}) \text{ or } (i > j \text{ and } i > \bar{j})\Big\} \subseteq S$, *represented by solid vertices in Figure 3.*

(2) $\Big\{(ij) \mid (i > j \text{ and } i < \bar{j}) \text{ or } (i < j \text{ and } i > \bar{j})\Big\} \cap S = \emptyset$, *represented by dotted vertices in Figure 3.*

(3) For all $1 \le i \le \frac{m-1}{2}$ there are two mutually exclusive cases, (a) and (b), represented by the two vertices x_i and the two vertices $\overline{x_i}$, respectively, in Figure 3.

(a) $\{(ii), (\bar{i}\bar{i})\} \subseteq S$ *and* $\{(i\bar{i}), (\bar{i}i)\} \cap S = \emptyset$,

(b) $\{(i\bar{i}), (\bar{i}i)\} \subseteq S$ *and* $\{(ii), (\bar{i}\bar{i})\} \cap S = \emptyset$.

These choices result in $2^{\frac{m-1}{2}}$ valid sets.

Proof. As indicated in Figure 3, for each tuple (ij) we call i the *row* of (ij). For S to be valid, according to Def. 6.8(3), S must contain $|i - \bar{i}|$ elements in row i and according to Def. 6.8(2), $\tau(s) \notin S$ for all $s \in S$.

In particular, S must contain $m - 1$ elements in row 1. If $(11) \in S$, then $(1m) \notin S$, because $\tau(11) = (1m)$. Hence there are only two possibilities: **(a):** $\{(1j) \mid 1 \le j < m\} \subseteq S$ or **(b):** $\{(1j) \mid 1 < j \le m\} \subseteq S$. By symmetry, for row m we get **(a')**: $\{(mj) \mid 1 \le j < m\} \subseteq S$ or **(b')**: $\{(mj) \mid 1 < j \le m\} \subseteq S$. But since $\tau(1m) = (mm)$ and $\tau(m1) = (11)$, the fact $\tau(S) = \iota(S)$ implies that **(a)** iff **(b')** and that **(a')** iff **(b)**. We are left with the two possibilities $\big($**(a)** and **(b')**$\big)$ or $\big($**(a')** and **(b)**$\big)$.

Now consider row 2. We have $\tau(21) = (1, m-1) \in S$ and hence $(21) \notin S$. In the same manner we see $(2m) \notin S$. We are left to choose $m - 3$ elements from the $m - 2$ remaining elements in row 2. The same argument as for row 1 gives two possibilities: **(a):** $\{(2j) \mid 2 \le j < m-1\} \subseteq S$ or **(b):** $\{(2j) \mid 2 < j \le m-1\} \subseteq S$. Analogously for row $m - 1$ we have **(a'):** $\{(m-1, j) \mid 2 \le j < m-1\} \subseteq S$ or **(b'):** $\{(m-1, j) \mid 2 < j \le m-1\} \subseteq S$. With the same reasoning as for the rows 1 and m we get **(a)** iff **(b')** and that **(a')** iff **(b)**. Again we are left with the two possibilities $\big($**(a)** and **(b')**$\big)$ or $\big($**(a')** and **(b)**$\big)$.

Continuing these arguments we end up with $2^{\frac{m-1}{2}}$ possibilities. It is easy to see that each of these possibilities gives a valid set. $\square$

The following claim finishes the proof of Prop. 6.1(2).

Claim 6.11. *All nonzero triple labelings J have the same coefficient of $\mathcal{X}$ in $\mathsf{eval}_{\mathcal{H}}(AJ)$.*

Proof. Take two nonzero triple labelings J and J'. According to Prop. 6.9, both sets $J^{(1)}(V^{(3)})$ and ${J'}^{(1)}(V^{(3)})$ are valid sets. Because of Lem. 6.10, it suffices to consider only the case where $J^{(1)}(V^{(3)})$ and $J'^{(1)}(V^{(3)})$ differ by a single involution $\sigma\colon O_m \to O_m$, where for some fixed $1 \le i \le \frac{m-1}{2}$ we have $\sigma(ii) = (i\bar{i})$ and $\sigma(\bar{i}\bar{i}) = (\bar{i}i)$, and σ is constant on all other pairs.

We analyze the labels that are affected by σ. We only perform the analysis for one of the two symmetric cases, namely for $\{|ii\rangle, |\bar{i}\bar{i}\rangle\} \subseteq J^{(1)}(V^{(3)})$. Note that this implies

$$\big\{\big(|ii\rangle, |i\bar{i}\rangle, |\bar{i}i\rangle\big), \big(|\bar{i}\bar{i}\rangle, |\bar{i}i\rangle, |i\bar{i}\rangle\big)\big\} \subseteq J(V^{(3)}), \qquad (\diamondsuit)$$

according to Claim 6.6. We adapt the notation from (6.1) to our special situation and write $t_{000} := t_{\bar{i}\bar{i}\bar{i}}$, $t_{001} := t_{\bar{i}\bar{i}i}$, $\ldots$, $t_{111} := t_{iii}$. Using this notation, $(\diamondsuit)$ reads as follows: $\{t_{110}, t_{001}\} \subseteq J(V^{(3)})$. Using Claim 6.7 we get

$$\{t_{101}, t_{010}\} \subseteq J(V^{(2)}), \quad \{t_{011}, t_{100}\} \subseteq J(V^{(1)}).$$

Applying σ to $J^{(1)}(V^{(3)})$, we can use Claim 6.6 again to get

$$\big\{\big(|i\bar{i}\rangle, |\bar{i}\bar{i}\rangle, |\bar{i}i\rangle\big), \big(|\bar{i}i\rangle, |ii\rangle, |i\bar{i}\rangle\big)\big\} \subseteq J'(V^{(3)}).$$

Applying Claim 6.7 and using our short syntax, we get:

$$\{t_{100}, t_{011}\} \subseteq J'(V^{(3)}),$$
$$\{t_{001}, t_{110}\} \subseteq J'(V^{(2)}),$$
$$\{t_{010}, t_{101}\} \subseteq J'(V^{(1)}).$$

We see that exactly the same triples occur in $J(V(\mathcal{H}))$ as in $J'(V(\mathcal{H}))$. We focus now on $J^{(1)}$ and $J'^{(1)}$ and see that:

$$\{(ii), (\bar{i}\bar{i})\} \subseteq J^{(1)}(V^{(3)}) \text{ and } \{(i\bar{i}), (\bar{i}i)\} \subseteq J^{(1)}(V^{(2)})$$

and

$$\{(i\bar{i}), (\bar{i}i)\} \subseteq {J'}^{(1)}(V^{(3)}) \text{ and } \{(\bar{i}\bar{i}), (ii)\} \subseteq {J'}^{(1)}(V^{(2)}).$$

This gives exactly two switches of positions in $e^{(1)} = V^{(2)} \mathbin{\dot\cup} V^{(3)} \mathbin{\dot\cup} \{y^0\}$, hence

$$\mathsf{eval}_{e^{(1)}}(AJ) = (-1)^2 \mathsf{eval}_{e^{(1)}}(AJ') = \mathsf{eval}_{e^{(1)}}(AJ').$$

Analogously we can prove that $\mathsf{eval}_{e^{(k)}}(AJ) = \mathsf{eval}_{e^{(k)}}(AJ')$ for all $k \in \{2, 3\}$ and therefore $\mathsf{eval}_{\mathcal{H}}(AJ) = \mathsf{eval}_{\mathcal{H}}(AJ')$. $\square$

Prop. 6.1 is completely proved. The same proof gives a lower bound on the *s-rank* [5] of the matrix multiplication tensor. More specifically, consider for $\alpha \in (\mathbb{C}^\times)^{m \times m \times m}$:

$$\mathcal{M}_m^\alpha := \sum_{i,j,l=1}^{m} \alpha_{ijl} |ij\rangle \otimes |jl\rangle \otimes |li\rangle.$$

Corollary 6.12. *For the obstruction design $\mathcal{H}$ and the same marix A as in the proof of Prop. 6.1, we have $f_{\mathcal{H}}(A\mathcal{M}_m^\alpha) \neq 0$. Hence $\underline{R}(\mathcal{M}_m^\alpha) \ge \frac{3}{2}m^2 - 2$ for all $\alpha \in (\mathbb{C}^\times)^{m \times m \times m}$.*

Let $\mathrm{mult}_\lambda(\mathscr{V})$ denote the multiplicity of $\{\lambda\}$ in the GL_n-representation $\mathscr{V}$. By restriction of functions, we have

$$\mathrm{mult}_\lambda(\mathbb{C}[\overline{\mathsf{GL}_n^3 \mathcal{E}_n}]_d) \le \mathrm{mult}_\lambda(\mathbb{C}[\mathsf{GL}_n^3 \mathcal{E}_n]_d). \qquad (6.3)$$

Remark 6.13. We can prove $\mathrm{mult}_\lambda(\mathbb{C}[\mathsf{GL}_n^3 \mathcal{E}_n]_d) = 0$, which is stronger than Prop. 6.1(1), cf. [9, Prop. 8.3.1].

7. DETERMINANTAL COMPLEXITY

We now turn from the tensor scenario to the polynomial scenario. Our goal is to find polynomials in the vanishing ideal of $\mathsf{GL}_{n^2}\det_n$, cf. [12]. For $\lambda \vdash_{n^2} dn$, let $p_\lambda(d[n])$ denote the multiplicity of $\{\lambda\}$ in the *plethysm* $\mathsf{Sym}^d\mathsf{Sym}^n\mathbb{C}^{n^2}$. From [4, eq. (5.2.6)] we know that

$$\mathbb{C}[\mathsf{GL}_{n^2}\det_n]_{\geq 0} = \bigoplus_{d\geq 0}\bigoplus_{\lambda \vdash_{n^2} nd} \mathsf{sk}\big(\lambda; (n\times d)^2\big)\{\lambda^*\}, \qquad (7.1)$$

where $\mathsf{sk}\big(\lambda; (n\times d)^2\big)$ is the *symmetric Kronecker coefficients*, defined in [4]. A sufficient criterion for the existence of a HWV of weight λ^* in the vanishing ideal $I(\mathsf{GL}_{n^2}\det_n)$ is given by

$$p_\lambda(d[n]) > \mathsf{sk}\big(n\times d; (\lambda)^2\big), \qquad (7.2)$$

since $\mathrm{mult}_{\lambda^*}(I(\mathsf{GL}_{n^2}\det_n)) \overset{(6.3),(7.1)}{\geq} p_\lambda(d[n]) - \mathsf{sk}\big(n\times d; (\lambda)^2\big)$.

Here are two examples of partitions satisfying (7.2), found by a computer: $(13,13,2,2,2,2,2) \vdash_7 36$ in degree $\frac{36}{3} = 12$ and $(15,5,5,5,5,5,5) \vdash_7 45$ in degree $\frac{45}{3} = 15$. An abundance of other partitions satisfying (7.2) is given in [9, Appendix].

The fact that a partition with 7 rows occurs in the vanishing ideal $I(\mathsf{GL}_9\det_3)_{12} \subseteq \mathsf{Sym}^{12}\mathsf{Sym}^3(\mathbb{C}^9)^*$ implies that the same partition occurs in the intersection $I(\mathsf{GL}_9\det_3) \cap \mathsf{Sym}^{12}\mathsf{Sym}^3(\mathbb{C}^7)^*$, see the inheritance theorems in [4]. Hence we get $f \notin \overline{\mathsf{GL}_9\det_3}$ for Zariski almost all polynomials $f \in \mathsf{Sym}^3(\mathbb{C}^7)^*$. Note that an explicit construction and evaluation of the HWVs in $\mathsf{Sym}^d\mathsf{Sym}^n\mathbb{C}^\ell$ directly gives lower bounds for docc for specific f.

8. REFERENCES

[1] S. Brunetti, A. Del Lungo, and Y. Gerard, *On the computational complexity of reconstructing three-dimensional lattice sets from their two-dimensional X-rays*, Proceedings of the Workshop on Discrete Tomography: Algorithms and Applications (Certosa di Pontignano, 2000), vol. 339, 2001, pp. 59–73.

[2] Peter Bürgisser, Michael Clausen, and M. Amin Shokrollahi, *Algebraic complexity theory*, Grundlehren der Mathematischen Wissenschaften, vol. 315, Springer-Verlag, Berlin, 1997.

[3] Peter Bürgisser and Christian Ikenmeyer, *Geometric complexity theory and tensor rank*, Proceedings 43rd Annual ACM Symposium on Theory of Computing 2011 (2011), 509–518.

[4] Peter Bürgisser, J.M. Landsberg, Laurent Manivel, and Jerzy Weyman, *An overview of mathematical issues arising in the Geometric complexity theory approach to VP v.s. VNP*, SIAM J. Comput. **40(4)** (2011), 1179–1209.

[5] Henry Cohn and Christopher Umans, *Fast matrix multiplication using coherent configurations*, (2012), arXiv:1207.6528 [math.NA].

[6] Igor Dolgachev, *Lectures on invariant theory*, London Mathematical Society Lecture Note Series, vol. 296, Cambridge University Press, Cambridge, 2003.

[7] P. C. Fishburn, J. C. Lagarias, J. A. Reeds, and L. A. Shepp, *Sets uniquely determined by projections on axes. II. Discrete case*, Discrete Math. **91** (1991), no. 2, 149–159.

[8] Roe Goodman and Nolan R. Wallach, *Symmetry, representations, and invariants*, Graduate Texts in Mathematics, vol. 255, Springer, Dordrecht, 2009.

[9] Christian Ikenmeyer, *Geometric complexity theory, tensor rank, and Littlewood-Richardson coefficients*, Ph.D. thesis, Institute of Mathematics, University of Paderborn, 2012, Online available at `http://math-www.upb.de/agpb/work/ikenmeyer_thesis.pdf`.

[10] Hanspeter Kraft, *Geometrische Methoden in der Invariantentheorie*, Friedr. Vieweg und Sohn Verlagsgesellschaft, Braunschweig, 1985.

[11] J. M. Landsberg, *The border rank of the multiplication of 2 x 2 matrices is seven*, J. Amer. Math. Soc **19** (2005), 447–459.

[12] J. M. Landsberg, Laurent Manivel, and Nicolas Ressayre, *Hypersurfaces with degenerate duals and the geometric complexity theory program*, (2010), arXiv:1004.4802.

[13] J. M. Landsberg and Giorgio Ottaviani, *New lower bounds for the border rank of matrix multiplication*, (2011), arXiv:1112.6007v2 [cs.CC].

[14] Joseph M. Landsberg and Laurent Manivel, *On the ideals of secant varieties of Segre varieties*, Found. Comput. Math. **4** (2004), no. 4, 397–422.

[15] Thomas Lickteig, *A note on border rank*, Inf. Process. Lett. **18** (1984), no. 3, 173–178.

[16] Ketan D. Mulmuley and Milind Sohoni, *Geometric complexity theory. I. An approach to the P vs. NP and related problems*, SIAM J. Comput. **31** (2001), no. 2, 496–526 (electronic).

[17] ———, *Geometric complexity theory. II. Towards explicit obstructions for embeddings among class varieties*, SIAM J. Comput. **38** (2008), no. 3, 1175–1206.

[18] Jeffrey B. Remmel, *A formula for the Kronecker products of Schur functions of hook shapes*, J. Algebra **120** (1989), no. 1, 100–118.

[19] Volker Strassen, *Vermeidung von Divisionen*, J. Reine Angew. Math. **264** (1973), 184–202.

[20] Volker Strassen, *Rank and optimal computation of generic tensors*, Lin. Alg. Appl. **52** (1983), 645–685.

[21] Volker Strassen, *Relative bilinear complexity and matrix multiplication*, J. Reine Angew. Math. **375/376** (1987), 406–443.

[22] Adolfo Torres-Cházaro and Ernesto Vallejo, *Sets of uniqueness and minimal matrices*, J. Algebra **208** (1998), no. 2, 444–451.

[23] L. G. Valiant, *Completeness classes in algebra*, Conference Record of the Eleventh Annual ACM Symposium on Theory of Computing (Atlanta, Ga., 1979), ACM, New York, 1979, pp. 249–261.

APPENDIX

We provide here the remaining proofs.

Proof of Lemma 5.5. We choose an arbitrary triple list $\xi \in \left((\mathbb{C}^n)^3\right)^d$ and show the vanishing of the contraction $\langle\widehat{\lambda}|\pi\mathcal{P}_d|\xi\rangle$, where $\pi \in \mathsf{S}_d^3$ corresponds to $\mathcal{H}$. According to Cor. 5.6 we have

$$\langle\widehat{\lambda}|\pi\mathcal{P}_d|\xi\rangle = \frac{1}{d!}\sum_{\vartheta\in\mathsf{S}_d\xi}\mathsf{eval}_{\mathcal{H}}(\vartheta).$$

We are going to see that switching the label of y and y' pairs summands that add up to zero: Let $\tau\colon V(\mathcal{H})\to V(\mathcal{H})$ denote the transposition switching y and y'. From a labeling $\vartheta\colon V(\mathcal{H})\to(\mathbb{C}^n)^3$ we get a new labeling $\vartheta\circ\tau$ by composition of maps. We show that $\mathsf{eval}_{\mathcal{H}}(\vartheta) = -\mathsf{eval}_{\mathcal{H}}(\vartheta\circ\tau)$. Let Δ denote the set of the three hyperedges containing both y and y'. For a hyperedge $e\in E^{(k)}$ we write $\vartheta_{|_e} := \vartheta^{(k)}{}_{|_e}$. Let

$$\alpha := \prod_{e\notin\Delta}\mathsf{eval}_e(\vartheta_{|_e}) = \prod_{e\notin\Delta}\mathsf{eval}_e((\vartheta\circ\tau)_{|_e}).$$

We calculate

$$\begin{aligned}\mathsf{eval}_{\mathcal{H}}(\vartheta)+\mathsf{eval}_{\mathcal{H}}(\vartheta\circ\tau) &= \alpha\prod_{e\in\Delta}\left(\mathsf{eval}_e(\vartheta_{|_e})+\mathsf{eval}_e((\vartheta\circ\tau)_{|_e})\right)\\ &= \alpha\prod_{e\in\Delta}\left(\mathsf{eval}_e(\vartheta_{|_e})+(-1)^3\mathsf{eval}_e(\vartheta_{|_e})\right)\\ &= 0. \qquad \square\end{aligned}$$

Proof of Claim in Example 5.10. According to Theorem 5.7, it suffices to show the result for HWVs $f_{\mathcal{H}}$ which correspond to obstruction designs $\mathcal{H}$ of type $\lambda = (n{\times}l)^3$. Let $A\in(\mathbb{C}^{n\times n})^3$ be arbitrary and let $|w_i^{(k)}\rangle := A^{(k)}|i\rangle$ for $1\le k\le 3$. We have $A\mathcal{E}_n = \sum_{i=1}^n |w_i^{(1)}\rangle\otimes|w_i^{(2)}\rangle\otimes|w_i^{(3)}\rangle \in \bigotimes^3\mathbb{C}^n$. Evaluating $f(A\mathcal{E}_n)$ as described in (5.3) yields

$$f(A\mathcal{E}_n) = \sum_{J\in\{1,\dots,n\}^{nl}}\mathsf{eval}_{\mathcal{H}}(w_{J_1},w_{J_2},\dots,w_{J_{nl}}) \qquad (*)$$

To get a nonzero summand it is necessary that for each hyperedge $e\subseteq\{1,\dots,nl\}$ we have that $J\colon\{1,\dots,nl\}\to\{1,\dots,n\}$ restricted to e is injective. In our special case, since $|e| = n$ for all hyperedges e, this means that J restricted to e is a bijection. Let Ψ denote the set of mappings $J\colon\{1,\dots,nl\}\to\{1,\dots,n\}$ which are bijective upon restriction to each hyperedge e. We define the involution $\varphi\colon\Psi\to\Psi$ by composing $\varphi(J) := (1\ 2)\circ J$, where $(1\ 2)\in\mathsf{S}_n$ switches 1 and 2. Clearly, $J\in\Psi$ iff $\varphi(J)\in\Psi$. Moreover,

$$\mathsf{eval}_{\mathcal{H}}(w_{J_1},\dots,w_{J_{nl}}) = (-1)^{3l}\mathsf{eval}_{\mathcal{H}}(w_{\varphi(J)_1},\dots,w_{\varphi(J)_{nl}}),$$

because the switch of 1 and 2 causes a sign change in each of the $3l$ hyperedges. Since $3l$ is an odd number, we paired summands in $(*)$ that add up to zero and hence the sum $f(A\mathcal{E}_n)$ evaluates to zero. $\square$

In the text it was claimed that $\mathcal{H}$ from Figure 2 is the only obstruction design of type $\lambda = ((2\kappa+1)\ulcorner(\kappa+1))^3$. We show here how this can be derived from general principles.

A set $r\in\mathscr{R}_\lambda$ is called *additive*, if there exist three real-valued functions $f^{(k)}\colon\mathbb{N}\to\mathbb{R}$ such that for all i,j,l

$$(i,j,l)\in r \Leftrightarrow f^{(1)}(i)+f^{(2)}(j)+f^{(3)}(l)\ge 0. \qquad (\dagger)$$

In [7, Thms. 1 and 2] it is shown that if $r\in\mathscr{R}_\lambda$ is additive, then $|\mathscr{R}_\lambda| = 1$.

The set $r := \{(i,1,1)\mid 1\le i\le\kappa+1\}\cup\{(1,j,1)\mid 1\le j\le\kappa+1\}\cup\{(1,1,l)\mid 1\le l\le\kappa+1\}\in\mathscr{R}_\lambda$ is easily seen to correspond to $\mathcal{H}$, cf. Subsec. 5.2.3. For all $1\le k\le 3$ we define $f^{(k)}\colon\mathbb{N}\to\mathbb{R}$ via $f^{(k)}(1) = 1$ and $f^{(k)}(i) = -1$ for all $i>1$. It is easy to check that $(\dagger)$ is satisfied and hence r is additive. We conclude $|\mathscr{R}_\lambda| = 1$ and hence there is exactly one obstruction design of type λ.